

ALLOCATION OF DIVISIBLE GOODS UNDER LEXICOGRAPHIC PREFERENCES

LEONARD J. SCHULMAN AND VIJAY V. VAZIRANI

ABSTRACT. We present a simple and natural non-pricing mechanism for allocating divisible goods among strategic agents having lexicographic preferences. Our mechanism has favorable properties of incentive compatibility (strategy-proofness), Pareto efficiency, envy-freeness, and time efficiency.

1. INTRODUCTION

The study of principled ways of allocating divisible goods among agents has long been a central topic in mathematical economics. The method of choice that emerged from this study, the Arrow-Debreu market model [1], provides a powerful approach based on pricing and leads to the fundamental welfare theorems. However, these market-based methods have limitations when agents are assumed to be strategic, e.g., these methods are not incentive compatible. Issues of the latter kind have been studied within the area of mechanism design for the last four decades, and have played a large role in the last decade in algorithmic game theory [18].

In this paper our primary focus is a particular simple and natural non-pricing mechanism, the *Synchronized Greedy (SG) mechanism*, for allocating divisible goods. The SG mechanism generalizes a mechanism introduced by Crès and Moulin [6] in the context of a job scheduling problem, and studied further by Bogomolnaia and Moulin [5] for the allocation of indivisible goods. For the setting defined below, we show that SG has favorable efficiency, incentive compatibility, and fairness properties. Our setting assumes that each agent has a lexicographic preference relation over goods. We note that this preference relation is *rational* in the sense that it is complete and transitive. It does not, on the other hand, satisfy the continuity condition that preferences between allocations are preserved under limits; a rational preference relation that also satisfies this continuity condition is known to be representable by a utility function, see [15]. However, the simplicity of the SG mechanism suggests the possibility of obtaining related mechanisms that achieve approximate versions of the above properties, when agents' preferences are representable by utility functions.

In detail, the allocation problem we consider is this. There are m distinct divisible goods which need to be allocated among n agents. Good j ($1 \leq j \leq m$) is available in the amount $q_j > 0$, and agent i ($1 \leq i \leq n$) is to receive a specified total of $r_i > 0$ across all goods; the parameters satisfy $\sum_j q_j = \sum_i r_i$. An *allocation* of goods is a list of numbers $a_{ij} \geq 0$, with $\sum_j a_{ij} = r_i$ and $\sum_i a_{ij} = q_j$, indicating that agent i receives quantity a_{ij} of good j . The vector $a_{i*} = (a_{i1}, \dots, a_{im})$ is referred to as agent i 's (share of the) allocation. Each agent i has a *preference list*, which is a permutation π_i of the goods;

Date: March 8, 2013.

$(a_{i\pi_i(1)}, \dots, a_{i\pi_i(m)})$ is agent i 's *sorted allocation*. Agent i 's preference among allocations is induced by *lexicographic order*. That is to say, agent i *lexicographic-prefers* a_{i*} to b_{i*} , $a_{i*} >_i b_{i*}$, if the leftmost nonzero coordinate of $(a_{i\pi_i(1)}, \dots, a_{i\pi_i(m)}) - (b_{i\pi_i(1)}, \dots, b_{i\pi_i(m)})$ is positive. Furthermore, we will say that agent i *majorization-prefers* a_{i*} to b_{i*} if

$$\text{for all } k = 1, \dots, m : \sum_{\ell=1}^k a_{i\pi_i(\ell)} \geq b_{i\pi_i(\ell)},$$

with at least one of the inequalities being strict. Observe that “lexicographic-prefers” is a complete preference relation without indifference contours (since it is antisymmetric for distinct allocation shares), and that “majorization-prefers” is an incomplete preference relation; moreover the lexicographic order is a refinement of the majorization order, i.e., majorization-prefers implies lexicographic-prefers. The phrase “agent i weakly X-prefers” will be used to include the possibility that agent i 's share is identical in the two allocations.

The SG mechanism has the following attributes w.r.t. the relation lexicographic-prefers.

- (1) The allocation produced by the SG mechanism in response to truthful bids is Pareto efficient.
- (2) If all r_i 's are equal, the allocation produced by the SG mechanism in response to truthful bids is envy-free in the following sense: each agent weakly majorization-prefers her allocation to that of any other agent.
- (3) Incentive compatibility: The SG mechanism is strategy-proof if $\min_j q_j \geq \max_i r_i$. We give a counterexample in the absence of this condition.
- (4) More generally: the SG mechanism is group strategy-proof against coalitions of ℓ agents if $\min_j q_j \geq \max_{S: |S|=\ell} \sum_{i \in S} r_i$. Again, we give a counterexample in the absence of this condition.
- (5) The running time to implement the SG mechanism is $\tilde{O}(mn)$.
- (6) An appropriate extension of the SG mechanism characterizes all Pareto efficient allocations. (However, in general, the extension does not possess the rest of the properties listed above.)

The SG mechanism is deterministic and treats all agents symmetrically.

It is also interesting to consider whether a mechanism is equitable—minimizing, in some measurable sense, the disparity in the welfare of the players. In spite of being deterministic and treating all agents symmetrically, the SG mechanism is not particularly equitable, except regarding the allocation of each agent's most preferred good. We provide an example showing that even the allocations of each agent's two most preferred goods may be quite inequitable. However, we describe a time-efficient algorithm that, for any given $1 \leq k \leq m$, equitably allocates the top k goods for each agent. We further define the notion of a *lexicographically most equitable* allocation and give a time-efficient algorithm to find one.

Since most of our paper deals with the relation “lexicographic-prefers”, we subsequently abbreviate it to “prefers”.

1.1. Literature. There has been considerable work on the strategy-proof allocation of divisible goods in Arrow-Debreu economies, starting with the seminal work of Hurwicz [10], e.g., see [7, 12, 20, 21, 22, 24]. Most of these results are negative, among the recent ones being Zhou's result showing that in a 2-agent, n -good pure exchange economy, there

can be no allocation mechanism that is efficient, non-dictatorial (i.e., both agents must receive non-zero allocations) and strategy-proof [24].

The paper that is most closely related to our work is that of Bogomolnaia and Moulin [5]. In their setting there are n agents and n indivisible goods, each agent having a total preference ordering over the goods; the desired outcome is a matching of goods with agents. A straightforward mechanism for allocating one good to each agent is *random priority (RP)*: pick a uniformly random permutation of the agents and ask each agent in turn to select a good among those left. It is easy to see that this mechanism is *ex post efficient*, i.e., the allocation it produces can be represented as a probability distribution over Pareto efficient deterministic allocations, and it is strategy-proof. However, it is not ex ante efficient. A random allocation is said to be *ex ante efficient* if for any profile of von Neumann-Morgenstern utilities that are consistent with the preferences of agents, the expected utility vector is Pareto efficient. It is easy to see that ex ante efficiency implies ex post efficiency.

Solving a conjecture of Gale [8], Zhou [23] showed that no strategy-proof mechanism that elicits von Neumann-Morgenstern utilities and achieves Pareto efficiency can find a “fair” solution even in the weak sense of equal treatment of equals. He further showed that the solution found by RP may not be efficient if agents are endowed with utilities that are consistent with their preferences. Hence, ex ante efficiency had to be sacrificed, if strategy-proofness and fairness were desired.

In the face of these choices, the work of Bogomolnaia and Moulin gave the notion of *ordinal efficiency* that is intermediate between ex post and ex ante efficiency; an allocation a is ordinally efficient if there is no other allocation b such that every agent majorization-prefers b to a . They went on to show that the mechanism called *probabilistic serial (PS)*, introduced in Crès and Moulin [6], yields an ordinally efficient allocation. Further they show that PS is envy-free and weakly strategy-proof, defined appropriately for the partial order “majorization-prefers”. Finally, Bogomolnaia and Moulin define an extension of PS by introducing different “eating rates” and show that this set of mechanisms characterizes the set of all ordinally efficient allocations.

Katta and Sethuraman [13] generalize the setting of Bogomolnaia and Moulin to the “full domain”, i.e., agents may be indifferent between pairs of goods. Thus, each agent partitions the goods by equality and defines a total order on the equivalence classes of her partition (the agent is equally happy with any good received from an equivalence class). For this setting, they give a randomized mechanism that is a generalization (different from ours) of PS and achieves the same game-theoretic properties as PS.

A mechanism that probabilistically allocates indivisible goods can also be viewed as one that fractionally allocates divisible goods. Under the latter interpretation, the SG mechanism is equivalent to PS for the case that $m = n$, and the quantity of each good and the requirement of each agent is one unit. An important difference is that Bogomolnaia and Moulin analyze PS under an incomplete preference relation (majorization) in which “most” allocation shares are incomparable; whereas we analyze SG under a complete preference relation (lexicographic) that is a refinement of majorization. The statement that a mechanism’s allocation is Pareto optimal w.r.t. lexicographic preferences is considerably stronger than the same statement w.r.t. majorization preferences, because each agent’s share is dominated by more alternative shares in the lexicographic order, than it is in the majorization order; so, fewer allocations are Pareto optimal in the lexicographic than

in the majorization order. Our results should be viewed therefore as demonstrating that the PS mechanism and its natural generalization, SG, have far stronger game-theoretic properties than even envisioned in [5].

Another setting in which incentive compatible, efficient mechanisms for allocating divisible goods have been studied is that of Leontief utilities (after imposing additional rules). Nicolo [17] gave a mechanism for the case of two goods and two agents; however, he could not generalize to an arbitrary number of goods and agents. This was achieved by Ghodsi et al. [9] under their *non-wasteful* rule. In general, under Leontief utilities, some of the resources may be redundant and will not improve any agent's welfare. The non-wasteful rule requires these excess resources to not be assigned to any agent and be removed (for possible use outside the mechanism). This removes the possibility of strategic manipulation, and in fact Ghodsi et al. gave an incentive compatible, efficient mechanism. A substantial generalization of this result was achieved by Li and Xue [14]; their mechanism is group strategy-proof and also satisfies fairness conditions.

Finally, we remark only that the problem of allocating a *single* divisible good among multiple agents with known privileges is considerably different; the principal issue studied in that problem is how to make the division in a manner that is fair w.r.t. the given privileges. This is known as the bankruptcy problem and has a long history, e.g., see [19, 2]. Despite an interesting resemblance between the PS mechanism and some of the mechanisms used in the solutions of that problem [11], the issues at stake in the bankruptcy literature are distinct from those in our paper and its predecessors.

2. THE SYNCHRONIZED GREEDY MECHANISM

The mechanism is simple. Each agent i submits a preference list σ_i , which is a permutation of the goods; the submitted list may or may not agree with his true preference list, called π_i . A simple, representative case, to keep in mind while reading the mechanism, is that of $m = n$ and all $q_j = r_i = 1$.

The mechanism simulates the following physical process. Considering good j as a “liquid”, and each agent as a receptacle of capacity r_i , the mechanism starts out at time 0 by (for all i in parallel) pouring good $\sigma_i(1)$ into receptacle i at rate r_i units of liquid per unit time. A good may be simultaneously poured into several receptacles. This continues until the first good is exhausted. Then instantaneously all agents whose favorite good ran out, switch to the next available good on their preference list, and begin receiving it at the same rate r_i . In general, whenever a good is exhausted, all agents who were in the process of receiving it, switch to the good which is highest in their preference permutation and has not yet been exhausted.

Each agent i is, at any time, receiving exactly one type of good, at rate r_i . At time 1 all agents simultaneously complete their full allocation.

This continuous process can easily be converted into an $\tilde{O}(mn)$ -time discrete algorithm: maintain a priority queue of goods, keyed by termination times.

3. PROPERTIES OF THE SYNCHRONIZED GREEDY MECHANISM

3.1. Pareto Optimality. Let a_{ij}^σ be the allocation created by the SG mechanism in response to bids σ declared by the agents. Let π denote the bids corresponding to the true preferences π_i .

Theorem 1. *The allocation produced by the SG mechanism in response to truthful bids is Pareto efficient: For all $a \neq a^\pi$ there is an i such that $a_{i*}^\pi >_i a_{i*}$.*

Proof. For a collection of bids σ let T_j^σ be the time at which good j is exhausted if the mechanism is run with bids σ . In particular T_j^π is the time at which good j is exhausted if the mechanism is run with the true preferences π . Let τ_1 be the first time at which any good is exhausted in π and let τ_k , $k \geq 2$, be the least time $t > \tau_{k-1}$ at which some good is exhausted. Let R_k be the set of goods which are exhausted at time τ_k .

By assumption $a \neq a^\pi$. For each j let $C_j^+ = \{i : a_{ij} < a_{ij}^\pi\}$ and let $C_j^- = \{i : a_{ij} > a_{ij}^\pi\}$. Let $J = \{j : C_j^+ \neq \emptyset\} = \{j : C_j^- \neq \emptyset\}$ be the set of goods that are allocated differently in a than in a^π .

Let k be least such that $R_k \cap J \neq \emptyset$; that is, τ_k is the first time at which a good of J is exhausted. Fix $j \in R_k \cap J$ and let $i \in C_j^+$. Then in order that $a_{i*} \geq_i a_{i*}^\pi$, there must be a j' , $\pi_i^{-1}(j') < \pi_i^{-1}(j)$, $a_{ij'} > a_{ij'}^\pi$. Then $i \in C_{j'}^-$ and so $j' \in J$. Moreover, since $a_{ij}^\pi > 0$, good j was available at the time that agent i requested it, which can only be after time $T_{j'}^\pi$, so $T_{j'}^\pi < T_j^\pi = \tau_k$. Letting k' be such that $\tau_{k'} = T_{j'}^\pi$, we have that $k' < k$ and $j' \in R_{k'} \cap J$, a contradiction. \square

3.2. Strategy-Proofness. A mechanism is said to be *strategy-proof* if no agent can obtain a strictly improved allocation by lying, provided the rest of the agents submit truthful bids.

Theorem 2. *The SG mechanism is strategy-proof if $\min q_j \geq \max r_i$.*

This is a special case of Theorem 7/Corollary 8. However for convenience we provide a stand-alone proof, since not all the complications of Theorem 7 arise.

Proof. Without loss of generality focus on agent 1. We need to show that for any bid σ_1 (with $\sigma = (\sigma_1, \pi_2, \dots, \pi_n)$), $a_{1*}^\sigma \leq_1 a_{1*}^\pi$. The theorem is trivial if $a^\sigma = a^\pi$.

The theorem is also trivial if agent 1, bidding truthfully, receives only his top choice. So we may suppose that agent 1 does not receive the entire allocation of any one good.

We may also suppose that if $a_{1j}^\sigma = 0$ and $a_{1j'}^\sigma > 0$, then $\sigma_1(j) > \sigma_1(j')$. In other words, all the requests in σ_1 that come up empty may as well be deferred to the end.

Let $\pi_1^{-1}(j)$ be the s such that $\pi_1(s) = j$. Let $G(j) = \{j' : \pi_1^{-1}(j') \leq \pi_1^{-1}(j) \text{ and } a_1^\pi(j') > 0\}$.

Say that agent 1 *sacrifices* good j in σ if:

- (1) $a_{1j}^\pi > 0$,
- (2) $\sigma_1(j) > |G(j)|$, and
- (3) $\pi_1^{-1}(j) < \pi_1^{-1}(j')$ if j' also satisfies (1),(2).

That is to say, j is the most-preferred good which agent 1 receives a positive quantity in π , but requests later in σ than in π .

Agent 1 must sacrifice some good, call it B , since otherwise the allocation will not change. See Figure 1.

Lemma 3. *If D is a good and $T_D^\pi < T_B^\pi$, then $T_D^\sigma \leq T_D^\pi$.*

Proof. Supposing the contrary, let D be a counterexample minimizing T_D^π . Since $T_D^\pi < T_B^\pi$, $D \neq B$.

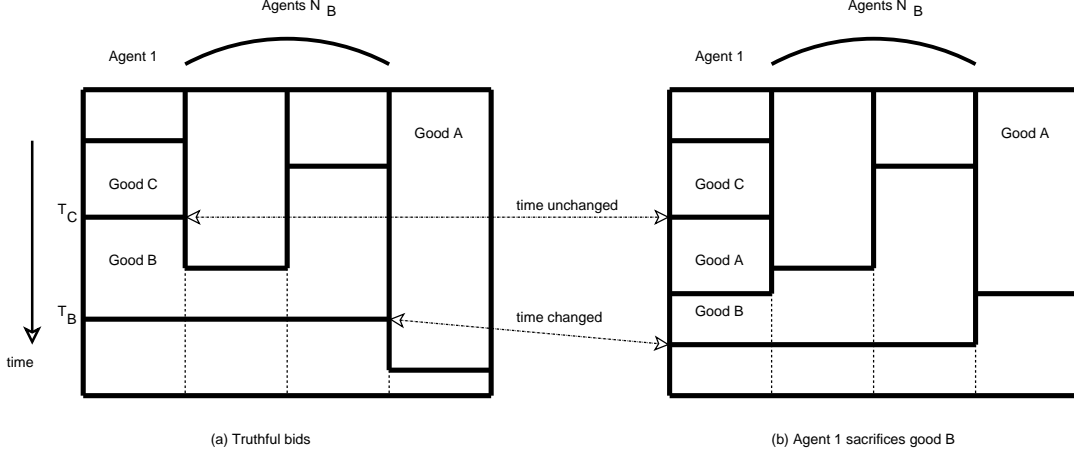


FIGURE 1. The mechanism with truthful vs. lying bids of Agent 1

Now let i be any agent (who may or may not be agent 1) for whom $a_{iD}^\pi > 0$. Due to the minimality of D , each of the goods j which i truthfully prefers to D , has $T_j^\sigma \leq T_j^\pi$. Therefore i requests D at a time in σ that is at least as soon as the time i requests it in π .

Since this holds for all i who received a positive allocation of D in π , the lemma follows. \square

Let N_B be the set of agents $i \neq 1$ for whom $a_{iB}^\pi > 0$. Due to the lemma, for each agent in $\{1\} \cup N_B$, the request time for B in σ is weakly earlier than it is in π . Now let C be the good such that $\pi_1^{-1}(C)$ is maximal subject to $\pi_1^{-1}(C) < \pi_1^{-1}(B)$ and $a_1^\pi(C) > 0$. Due to the lemma, all goods j' such that $\pi_1^{-1}(j') \leq \pi_1^{-1}(C)$ have $T_{j'}^\sigma \leq T_{j'}^\pi$. Next we show:

Proposition 4. *If $\pi_1^{-1}(j') \leq \pi_1^{-1}(C)$, then $a_{1j'}^\sigma = a_{1j'}^\pi$.*

Proof. Supposing the contrary, let $\pi_1^{-1}(j')$ be minimal such that $\pi_1^{-1}(j') \leq \pi_1^{-1}(C)$ and $a_{1j'}^\sigma \neq a_{1j'}^\pi$. There are two possibilities to consider.

- (a) $a_{1j'}^\sigma < a_{1j'}^\pi$. This is not possible because then $a_{1*}^\sigma < a_{1*}^\pi$.
- (b) $a_{1j'}^\sigma > a_{1j'}^\pi$. Note:

Lemma 5. *Let j_1, j_2 be such that $\pi_1^{-1}(j_1) \leq \pi_1^{-1}(B)$, $\pi_1^{-1}(j_2) \leq \pi_1^{-1}(B)$, $a_{1j_1}^\pi > 0$, and $\pi_1^{-1}(j_1) < \pi_1^{-1}(j_2)$. Then $\sigma_1^{-1}(j_1) < \sigma_1^{-1}(j_2)$.*

Proof. Consider the least j_1 that is part of a pair j_1, j_2 violating the lemma. Then j_1 satisfies conditions (1),(2) above, contradicting that B is the good sacrificed by agent 1. \square

It follows that $T_{j'}^\sigma \geq \sum_{j'': \pi_1^{-1}(j'') \leq \pi_1^{-1}(j')} a_{1j''}^\sigma$. Due to the minimality of j' , this means that if $a_{1j'}^\sigma > a_{1j'}^\pi$, then $T_{j'}^\sigma > T_{j'}^\pi$, contradicting our earlier conclusion. This completes demonstration of the Proposition. \square

A consequence of the Proposition is that $T_C^\sigma = T_C^\pi$.

Since agent 1 sacrifices B , his request time for B in σ is strictly greater than his request time for B in π .

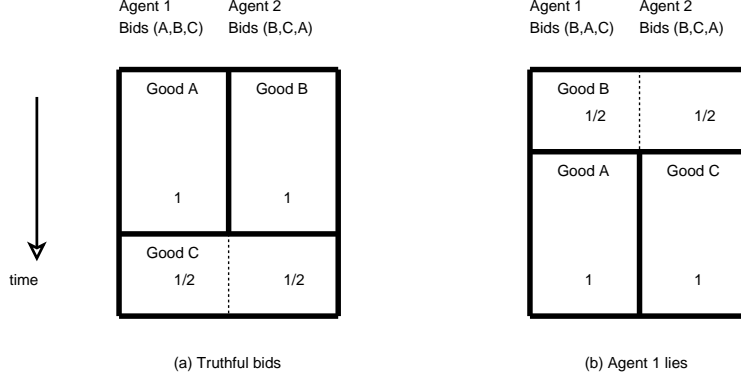


FIGURE 2. Failure of incentive compatibility without the hypothesis of Theorem 2

Recall that N_B is nonempty. At time T_B^π , the agents of N_B have received as least as much of B in σ as they have in π , and the latter is positive. On the other hand, at the same time T_B^π , agent 1 has received strictly less of B in σ than he has in π . In order for agent 1 to receive at least as much of B in σ as in π , he would have to receive all of B that is allocated after time T_B^π ; however, that is not possible, because the set of agents receiving B after T_B^π includes N_B . Thus $a_{1*}^\sigma <_1 a_{1*}^\pi$. \square

3.3. Necessity of a Hypothesis on $\{r_i\}, \{q_j\}$. We next provide an example in which strategy-proofness fails in the absence of the condition $\max r_i \leq \min q_j$. For convenience now let $r_1 \geq \dots \geq r_n$ and $q_1 \leq \dots \leq q_m$.

Example 6. Let $n = 2$ and $m = 3$. Let $r_1 = r_2 = 3/2$; label the goods A, B, C , let $q_A = q_B = q_C = 1$, and let the preference lists be $\pi_1 = (A, B, C)$, $\pi_2 = (B, C, A)$. If agent 1 bids truthfully he receives the sorted allocation $(1, 0, 1/2)$. If instead he bids (B, A, C) (while agent 2 bids truthfully), he receives the improved sorted allocation $(1, 1/2, 0)$. See Figure 2.

This example does not limit the theorem sharply, because it uses $r_1 = (3/2)q_1$ rather than r_1 arbitrarily close to q_1 . Jeremy Hurwitz has pointed out that one may construct similar examples with whenever $r_1 \geq q_1/(1 - q_2/\sum q_j)$; this would appear to be a tight bound.

3.4. Group Strategy-Proofness. A mechanism is *group strategy-proof* against a family F of subsets of agents if for every “coalition” $S \in F$, if all the agents outside S submit truthful bids, then the agents of S cannot obtain an improved allocation by lying, where by “improved allocation” we mean that no agent of S obtains a worse allocation and at least one obtains a strictly better allocation.

Theorem 7. *The SG mechanism is group strategy-proof against the family of subsets S for which $\min_j q_j \geq \sum_{i \in S} r_i$.*

Corollary 8. *The SG mechanism is group strategy-proof against coalitions of ℓ agents if $\min_j q_j \geq \max_{S: |S|=\ell} \sum_{i \in S} r_i$.*

Proof. Let S be a coalition as in the theorem statement. We need to show there is no list of bids for the agents in S such that all do at least as well as in π , and some do strictly better.

The structure of the proof is, as promised, the same as that of Theorem 2, but what makes this theorem more involved is that different agents in S can sacrifice different goods, and some of them may be better off due to the untruthful bids, as they can benefit from the results of others' lies; the proof needs to effectively "chase through" the complicated transfer of goods (relative to a^π), and show that *some* agent in the coalition is worse off than in π . Fortunately this "chasing" will not require any actual iteration.

If $\alpha_{i1}^\pi = 1$ for all $i \in S$, that is, with truthful bids these agents receive only their top choices, then none of them can be strictly rewarded by submitting a different bid.

Otherwise (i.e., if $\alpha_{i1}^\pi < 1$ for some $i \in S$), then thanks to the hypothesis, under the truthful bids π , every good has a positive allocation outside S .

We need to show that if some members of S submit untruthful bids, while the agents outside S bid truthfully, and if the resulting allocation is different than in π , then some agent in S does strictly worse than in π . Let σ be the alternate bids. (Note $\sigma_i = \pi_i$ for $i \notin S$.) We use the phrase " i is a willing participant in the coalition S " to mean that $i \in S$ and $a_{i*}^\sigma \geq_i a_{i*}^\pi$.

Without loss of generality we may suppose that every agent $i \in S$ bids untruthfully and that this has an effect, i.e., if the agent reverts to a truthful bid then the allocation is different than in σ .

Moreover, we may simplify the argument slightly by supposing that for each agent $i \in S$, if $a_{ij}^\sigma = 0$ and $a_{ij'}^\sigma > 0$, then $\sigma_i(j) > \sigma_i(j')$. In other words, all the requests that come up empty may as well be deferred to the end.

Let $\pi_i^{-1}(j)$ be the s such that $\pi_i(s) = j$. Let $G(i, j) = \{j' : \pi_i^{-1}(j') \leq \pi_i^{-1}(j) \text{ and } a_i^\pi(j') > 0\}$.

Say that agent i *sacrifices* good j in σ if:

- (1) $a_{ij}^\pi > 0$,
- (2) $\sigma_i(j) > |G(i, j)|$, and
- (3) $\pi_i^{-1}(j) < \pi_i^{-1}(j')$ if j' also satisfies (1),(2).

Some good must be sacrificed by some agent, since otherwise the allocation will not change. (However, while every agent in S is untruthful, not every $i \in S$ necessarily sacrifices a good; setting $\sigma_i(j) > \pi_i(j)$ might have an effect even if $a_i^\pi(j) = 0$ because of increased availability of j due to bidding changes of other agents.)

Of all the sacrificed goods let B be one for which T_B^π is minimal.

Lemma 9. *If D is a good and $T_D^\pi < T_B^\pi$, then $T_D^\sigma \leq T_D^\pi$.*

Proof. Supposing the contrary, let D be a counterexample minimizing T_D^π . By the minimality of B , D cannot be a sacrificed good.

Now let i be any agent (inside or outside of S) for whom $a_{iD}^\pi > 0$. Due to the minimality of D , each of the goods j which i truthfully prefers to D , has $a_j^\sigma \leq T_j^\pi$. Therefore i requests D at a time in σ that is at least as soon as the time i requests it in π .

Since this holds for all i who received a positive allocation of D in π , the lemma follows. \square

Let $O_B \subseteq S$ be the set of agents who sacrifice B , and let N_B be the set of agents i for whom $a_{iB}^\pi > 0$ but who do not sacrifice B . Due to the lemma, for each agent in $O_B \cup N_B$, the request time for B in σ is weakly earlier than it is in π . Now consider an agent $i \in O_B$. Let C be the good such that $\pi_i^{-1}(C)$ is maximal subject to $\pi_i^{-1}(C) < \pi_i^{-1}(B)$ and

$a_i^\pi(C) > 0$. Due to the lemma, all goods j' such that $\pi_i^{-1}(j') \leq \pi_i^{-1}(C)$ have $T_{j'}^\sigma \leq T_{j'}^\pi$. Next we show:

Proposition 10. *If $\pi_i^{-1}(j') \leq \pi_i^{-1}(C)$, then $a_{ij'}^\sigma = a_{ij'}^\pi$.*

Proof. Supposing the contrary, let $\pi_i^{-1}(j')$ be minimal such that $\pi_i^{-1}(j') \leq \pi_i^{-1}(C)$ and $a_{ij'}^\sigma \neq a_{ij'}^\pi$. There are two possibilities to consider.

- (a) $a_{ij'}^\sigma < a_{ij'}^\pi$. This is not possible because i is a willing participant in the coalition.
- (b) $a_{ij'}^\sigma > a_{ij'}^\pi$. Note:

Lemma 11. *Let j_1, j_2 be such that $\pi_i^{-1}(j_1) \leq \pi_i^{-1}(B)$, $\pi_i^{-1}(j_2) \leq \pi_i^{-1}(B)$, $a_{ij_1}^\pi > 0$, and $\pi_i^{-1}(j_1) < \pi_i^{-1}(j_2)$. Then $\sigma_i^{-1}(j_1) < \sigma_i^{-1}(j_2)$.*

Proof. Identical to the proof of Lemma 5 with agent i in place of agent 1. \square

It follows that $T_{j'}^\sigma \geq \sum_{j'': \pi_i^{-1}(j'') \leq \pi_i^{-1}(j')} a_{ij''}^\sigma$. Due to the minimality of j' , this means that if $a_{ij'}^\sigma > a_{ij'}^\pi$, then $T_{j'}^\sigma > T_{j'}^\pi$, contradicting our earlier conclusion. This completes demonstration of the Proposition. \square

A consequence of the Proposition is that $T_C^\sigma = T_C^\pi$.

Since agent i sacrifices B , his request time for B in σ is strictly greater than his request time for B in π .

Since we are in the case that every good has a positive allocation outside S , N_B is nonempty. At time T_B^π , the agents of N_B have received as least as much of B in σ as they have in π , and the latter is positive. On the other hand, at the same time T_B^π , the agents of O_B have received strictly less of B in σ than they have in π . In order for the agents of O_B to receive collectively at least as much of B in σ as in π , they would have to receive all of B that is allocated after time T_B^π ; however, that is not possible, because the set of agents receiving B after T_B^π includes N_B . Therefore there is some $i \in O_B$ for whom $a_{iB}^\sigma < a_{iB}^\pi$. This contradicts the requirement that i be a willing participant in the coalition S . \square

Example 12. Example 6, in which strategy-proofness failed absent the hypothesis of Theorem 2, can be extended in a straightforward manner to one in which the group strategy-proof property fails to hold absent the hypothesis of Corollary 8. Again use $m = 3$, but instead of two agents, use $n = 2\ell$ agents, the first half having the same preference order (A, B, C) as agent 1 in the earlier example, and the second half having the same preference order (B, C, A) as agent 2 in the earlier example. If all agents bid truthfully, then the first ℓ agents each receive the sorted allocation $(1, 0, 1/2)$; however if they lie and bid (B, A, C) , while the remainder bid truthfully, then each lying agent receives the improved sorted allocation $(1, 1/2, 0)$.

3.5. Envy-Freeness.

Theorem 13. *Suppose all r_i are equal. Under truthful bidding, every agent i weakly majorization-prefers his allocation a_{i*}^π to the allocation $a_{i'*}^\pi$ of any other agent i' .*

Proof. Fix any $1 \leq t \leq m$. We are to show that

$$\sum_{\ell=1}^k a_{i\pi_i(\ell)}^\pi \geq \sum_{\ell=1}^k a_{i'\pi_i(\ell)}^\pi.$$

Let $t = \max_{1 \leq \ell \leq k} T_{\pi_i(\ell)}^\pi$. Then $t / \sum q_j$ is the time that agent i stops receiving his top k goods. So $t = \sum_{\ell=1}^k a_{i\pi_i(\ell)}^\pi$. No other agent can receive any of these goods after time t , so $t \geq \sum_{\ell=1}^k a_{i'\pi_i(\ell)}^\pi$. \square

(If the r_i 's are not equal, the same statement applies to the relative allocations; see Sec. 5.)

4. CHARACTERIZING ALL PARETO EFFICIENT ALLOCATIONS

Bogomolnaia and Moulin [5] extended their mechanism by allowing players to receive goods at time-varying rates. Specifically, for each agent i there is a speed function η_i mapping the time interval $[0, 1]$ into the nonnegative reals, such that for all i

$$\int_0^1 \eta_i(t) dt = r_i.$$

Subject to these speeds, goods flow to agents in order of the preference lists they bid, just as before. They showed that this extension characterizes all ordinally efficient allocations. In this section, we obtain an analogous characterization of all Pareto efficient allocations by a similar extension of our mechanism. Specifically, we prove that for *any* Pareto efficient allocation of goods, there exist speeds such that the extended SG mechanism produces that allocation. We prove this after first noting that the extended SG mechanism always results in Pareto efficient allocations.

4.1. Pareto Efficiency.

Theorem 14. *Let η_i , $1 \leq i \leq n$, be speed functions. The allocation a^π produced by the extended SG mechanism with truthful bids and these speeds, is Pareto optimal.*

Proof. Let T_j^π be the time at which good j is exhausted with bids π and speeds η_i . The proof is then, word for word, the proof of Theorem 1. \square

4.2. Characterizing All Pareto Efficient Allocations. If the last result mirrored the First Welfare Theorem, the next mirrors the Second Welfare Theorem:

Theorem 15. *Let π be the collection of agent preference lists and let a be a Pareto-efficient allocation. There exist speed functions η_i , $1 \leq i \leq n$, such that $a = a^\pi$.*

Proof. Construction of the η_i is simple. Let a “partial allocation” be $\alpha_{ij} \geq 0$ such that $\sum_j \alpha_{ij} \leq r_i$ and $\sum_i \alpha_{ij} \leq q_j$.

Initialize $t = 0$ and initialize each agent i with the empty partial allocation $\alpha_{i1} = \dots = \alpha_{im} = 0$. Initialize also $c_j = q_j$ for all j . Then repeat the following until all $t = 1$.

Find an agent i for whom there is an ℓ such that $\alpha_{i\pi_i(\ell)} < a_{i\pi_i(\ell)}$, and such that for all $\ell' < \ell$, $c_{\pi_i(\ell')} = 0$. Set $\delta = (a_{i\pi_i(\ell)} - \alpha_{i\pi_i(\ell)}) / \sum q_j$. For $t < t' < t + \delta$, make the settings $\eta_i(t') = \sum q_j$ and, for $i' \neq i$, $\eta_{i'}(t') = 0$. Then increment $\alpha_{i\pi_i(\ell)}$ by $\delta \sum q_j$, and decrement $c_{\pi_i(\ell)}$ by the same amount. Finally, increment t by δ .

This process can only fail to complete if there comes a time t at which every agent i falls into one of the following sets S_1 and S_2 , and S_1 is nonempty:

- (1) $S_1 = \{i \text{ such that } \sum_j \alpha_{ij} < r_i, \text{ and the minimal } \ell \text{ for which } c_{\pi_i(\ell)} > 0 \text{ also satisfies } \alpha_{i\pi_i(\ell)} = a_{i\pi_i(\ell)}\}.$
- (2) $S_2 = \{i \text{ such that } \sum_j \alpha_{ij} = r_i\}.$

If there is such a t , then for each $i \in S_1$, define ℓ_i to be the ℓ identified in the definition of S_1 . Then for each $i \in S_1$, a small positive amount of the good $\pi_i(\ell_i)$ can be added to i 's current partial allocation; the new partial allocation, no matter how it is completed to an allocation, improves strictly on a for all $i \in S_1$, and is unchanged for $i \in S_2$. Therefore a is not Pareto efficient. \square

Examination of the above proof reveals:

Corollary 16. *There is a polynomial time algorithm for checking whether a given allocation is Pareto efficient.*

4.3. No Incentive Compatibility for the Variable Speeds Variant. We note that the synchrony imposed among agents by the SG mechanism is key to its incentive compatibility and envy-freeness properties (indeed, the properties hold even if the basic mechanism is extended with the *same* speed function for all agents). If different agents have different speed functions under the extended SG mechanism, Theorems 2 and 7, showing incentive compatibility, fail to hold. The argument breaks down as soon as it uses termination times, in Lemma 3. Below is a counter-example for strategy-proofness; a similar idea gives counter-examples for group strategy-proofness and envy-freeness.

Example 17. Assume $m = n = 4$ and that all $r_i = q_j = 1$. Let the speed function for agent 1 be 1 over the interval $[0, 1]$. The speeds of agents 2, 3, and 4 equal 1 over the interval $[0, 1/2]$, 0 over the interval $(1/2, 5/6]$, and 3 over the interval $(5/6, 1]$. The preference orders of agents 1 and 2 are $(1, 2, 3, 4)$, and the preference orders of agents 3 and 4 are $(2, 4, 3, 1)$. If all agents bid truthfully, agent 1 receives the sorted allocation $(1/2, 0, 1/2, 0)$. On the other hand, if agent 1 bids $(2, 1, 3, 4)$ while the rest bid truthfully, then agent 1 receives the better sorted allocation $(1/2, 1/3, 1/6, 0)$.

5. EQUITABLE ALLOCATIONS

Given an allocation a , let \bar{a} denote the *relative allocation*, where $\bar{a}_{ij} = a_{ij}/r_i$. For any k , $1 \leq k \leq m$, say that an allocation is *equitable w.r.t. agents' top k choices* if it belongs to

$$\operatorname{argmax}_a \min_i (\bar{a}_{i\pi_i(1)} + \dots + \bar{a}_{i\pi_i(k)}),$$

where the max is over all allocations a .

It is easy to see that the allocation produced by the SG mechanism is equitable for $k = 1$. However, as the following example illustrates, it is not equitable for $k = 2$, or larger values of k .

Example 18. Let $n = 2$, $m = 3$, $r_1 = r_2 = 1$, $q_1 = 1/2$, $q_2 = 5/6$, and $q_3 = 2/3$. Let the preference list of the first agent be $(1, 2, 3)$ and that of the second agent $(2, 3, 1)$. Then the SG mechanism gives sorted allocations of $(1/2, 1/6, 1/3)$ and $(2/3, 1/3, 0)$ respectively to the agents, so each receives $2/3$ of his total allocation from his top two choices. On the other hand, the sorted allocations $(1/2, 1/2, 0)$, $(1/3, 2/3, 0)$ are also feasible, and in this case each agent receives his entire allocation from his top two choices.

Next, we show that there is a polynomial-time algorithm which, given k , (r_i) , (q_j) and the list of agent preferences, obtains an allocation that is equitable w.r.t. agents' top k choices. In fact this allocation $x = (x_{ij})$ is the solution to the linear program given below,

together with t , the minimum over agents of the relative allocation from the agent's top k goods.

$$\begin{aligned}
 (5.1) \quad & \text{Maximize } t \\
 & \text{Such that } \forall i : t \leq \frac{1}{r_i} \left(\sum_{\ell=1}^k x_{i\pi_i(\ell)} \right) \\
 & \forall i : \sum_{j=1}^m x_{ij} = r_i \\
 & \forall j : \sum_{i=1}^n x_{ij} = q_j \\
 & \forall i \forall j : x_{ij} \geq 0
 \end{aligned}$$

Finally, let us define the notion of the *lexicographically most equitable allocation*, which intuitively is an allocation that simultaneously optimizes for each k , to the extent possible. For any allocation a , and each k , $1 \leq k \leq m$, define

$$(5.2) \quad \beta_k = \min_i (\bar{a}_{i\pi_i(1)} + \dots + \bar{a}_{i\pi_i(k)}).$$

Now, define a *lexicographically most equitable allocation* to be one that lexicographically maximizes $(\beta_1, \dots, \beta_m)$.

We now give a polynomial-time algorithm to find a lexicographically most equitable allocation—it involves solving m LPs derived from LP (5.1). The first LP simply computes β_1 by solving LP (5.1) for $k = 1$. Next, for each k , $2 \leq k \leq m$, add the following constraints to LP (5.1) and solve it to determine β_k :

$$\forall i, \forall 1 \leq h \leq k-1 : \frac{1}{r_i} \left(\sum_{\ell=1}^h x_{i\pi_i(\ell)} \right) \geq \beta_h.$$

Clearly, the last LP will yield a most equitable allocation.

Example 19. For the agents in Example 18, the lexicographically most equitable allocation is (given as a sorted allocation): $(1/2, 1/3, 1/6)$ for agent 1 and $(1/2, 1/2, 0)$ for agent 2. This is different from both the SG allocation and the allocation that is equitable w.r.t. agents' top 2 choices.

Although equitability seems to be a desirable property, it must be noted that an equitable allocation need not be even Pareto optimal:

Example 20. Let $n = 3$, $m = 4$, $r_1 = r_2 = r_3 = 2$, $q_1 = q_2 = 1$, $q_3 = q_4 = 2$. Let the preference lists be $\pi_1 = (1, 2, 3, 4)$, $\pi_2 = (3, 4, 1, 2)$, $\pi_3 = (4, 3, 1, 2)$. For any $0 \leq x \leq 1$, the following allocation is lexicographically most equitable, and even stronger, it simultaneously optimizes all β_k in Eqn. 5.2:

$$a_1 = (1, 1, 0, 0), a_2 = (0, 0, 1, x), a_3 = (0, 0, 1 - x, 1).$$

Yet this allocation is Pareto optimal only in the single case $x = 0$.

6. DISCUSSION

Our main open problem is the one mentioned in the Introduction, i.e., achieving approximate versions of the properties of the SG mechanism but when agents' preferences are representable by utility functions.

In [4], Bogomolnaia and Heo show that efficiency (under the majorization relation), envy-freeness, and a property they call bounded invariance characterize the PS mechanism. This leads to the question of appropriately characterizing the SG mechanism. Towards this end we ask if efficiency, under the more stringent lexicographic relation, and envy-freeness suffice. Clearly, a first step would be to characterize the PS mechanism in this manner. A mechanism is said to have the *bounded invariance property* if for any agent i and any good j , changing i 's preference order for goods she likes less than j does not change the amount (equivalently, probability) of good j each agent gets.

A natural open question concerns the existence of mechanisms to produce lexicographically most equitable allocations, having favorable algorithmic and game-theoretic properties (e.g., incentive compatibility).

Consider the generalization of our setting so agents may be indifferent between pairs of goods. Thus, each agent partitions the goods by equality and defines a total order on the equivalence classes of her partition (the agent is equally happy with any good received from an equivalence class). Preferences are again defined lexicographically over classes, i.e., by considering the total amount of goods received in each class. Is there a generalization of our mechanism to this setting?

Finally, consider the setting in which agents' preferences are random permutations of $1, \dots, m$. This leads to many interesting questions, e.g., what is the distribution of β_k as in Eqn. 5.2, for the allocation a given by the SG mechanism; and, what is the distribution of the maximized value of t as in LP (5.1), both for various values of k . Regarding the latter question, for $n = m$ and all $r_i = q_j = 1$, there is a correspondence in the case $k = 1$ with the collision statistics of random pointers, and so it is known that $t \rightarrow (\log \log n)/(\log n)$; for larger k there is a rough correspondence with the "power of two choices" literature [3, 16], suggesting likely asymptotics of $(\log k)/(\log \log n)$ for fixed k , although the correspondence between the problems is not close enough for us to state this with certainty.

7. ACKNOWLEDGMENTS

We are indebted to Hervé Moulin for generously sharing his deep understanding of this research domain. Thanks also to Jeremy Hurwitz for stimulating discussions and to Amin Saberi for pointing us to useful references.

Schulman was supported in part by the NSF. Vazirani was supported in part by NSF Grant CCF-0914732 and a Google Research Grant. Vazirani would like to thank the Social and Informational Sciences Laboratory at Caltech for their hospitality; this work was done while he was Distinguished SISL Visitor during 2011-12.

REFERENCES

- [1] K. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22:265–290, 1954.
- [2] R. Aumann and M. Maschler. Game theoretic analysis of a bankruptcy problem from the Talmud. *J. Economic Theory*, 36:195–213, 1985.

- [3] Y. Azar, A. Z. Broder, A. R. Karlin, and E. Upfal. Balanced allocations. *SIAM J. Comput.*, 29(1):180–200, 1999. (STOC 1994:593–602).
- [4] A. Bogomolnaia and E. J. Heo. Probabilistic assignment of objects: Characterizing the serial rule. *Journal of Economic Theory*, 2012. To appear.
- [5] A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100:295–328, 2001.
- [6] H. Crès and H. Moulin. Scheduling with opting out: improving upon random priority. *Operations Research*, 49(4):565–577, 2001.
- [7] P. Dasgupta, P. Hammond, and E. Maskin. The implementation of social choice rules. *Review of Economic Studies*, 46:153–170, 1979.
- [8] D. Gale. College course assignments and optimal lotteries. Mimeo, University of California, Berkeley, 1987.
- [9] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. Dominant resource fairness: fair allocation of multiple resource types. In *Proc. 8'th USENIX conference on networked systems design and implementation*, 2011.
- [10] L. Hurwicz. On informationally decentralized systems. In C. B. McGuire and R. Radner, editors, *Decision and Organization*, pages 297–336. North-Holland, Amsterdam, 1972.
- [11] M. Kaminski. Hydraulic rationing. *Mathematical Social Sciences*, 40:131–155, 2000.
- [12] M. Kato and S. Ohseto. Toward general impossibility theorems in pure exchange economies. *Social Choice and Welfare*, 19:659–664, 2002.
- [13] A. Katta and J. Sethuraman. A solution to the random assignment problem on the full preference domain. *Journal of Economic Theory*, 131:231–250, 2006.
- [14] J. Li and J. Xue. Egalitarian division under Leontief preferences. Manuscript, 2012.
- [15] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic theory*. Oxford University Press, 1995.
- [16] M. Mitzenmacher, A. Richa, and R. Sitaraman. The power of two random choices: A survey of techniques and results. In P. Pardalos, S. Rajasekaran, and J. Rolim, editors, *Handbook of Randomized Computing: Volume 1*, pages 255–312. To appear.
- [17] A. Nicolo. Efficiency and truthfulness with Leontief preferences. *Review of Economic Design*, 8(4):373–382, 2004.
- [18] N. Nisan. Introduction to mechanism design (for computer scientists). In N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, pages 209–241. Cambridge University Press, 2007.
- [19] B. O'Neill. A problem of rights arbitration from the Talmud. *Mathematical Social Sciences*, 2(4):345–371, 1982.
- [20] M. Satterthwaite and H. Sonnenschein. Strategy-proof allocation mechanisms at differentiable points. *Review of Economic Studies*, 48:587–597, 1981.
- [21] S. Serizawa. Inefficiency of strategy-proof rules for pure exchange economies. *Journal of Economic Theory*, 106:219–241, 2002.
- [22] S. Serizawa and J. Weymark. Efficient strategy-proof exchange and minimum consumption guarantees. *Journal of Economic Theory*, 109:246–263, 2003.
- [23] L. Zhou. On a conjecture by Gale about one-sided matching problems. *J. Econ. Theory*, 52:123–135, 1990.
- [24] L. Zhou. Inefficiency of strategy-proof allocation mechanisms in pure exchange economies. *Social Choice and Welfare*, 8(3):247–254, 1991.

CALTECH, MC305-16, PASADENA CA 91125, USA.

E-mail address: schulman@caltech.edu

COLLEGE OF COMPUTING, GEORGIA INSTITUTE OF TECHNOLOGY, ATLANTA GA 30332, USA.

E-mail address: vazirani@cc.gatech.edu